

Sparse Multinomial Logistic Regression via Approximate Message Passing

Evan Byrne and Philip Schniter*

Abstract—For the problem of multi-class linear classification and feature selection, we propose approximate message passing approaches to sparse multinomial logistic regression (MLR). First, we propose two algorithms based on the Hybrid Generalized Approximate Message Passing (HyGAMP) framework: one finds the maximum a posteriori (MAP) linear classifier and the other finds an approximation of the test-error-rate minimizing linear classifier. Then we design computationally simplified variants of these two algorithms. Next, we detail methods to tune the hyperparameters of their assumed statistical models using Stein’s unbiased risk estimate (SURE) and expectation-maximization (EM), respectively. Finally, using both synthetic and real-world datasets, we demonstrate improved error-rate and runtime performance relative to existing state-of-the-art approaches to sparse MLR.

Index Terms—Classification, feature selection, belief propagation, message passing.

I. INTRODUCTION

A. Objective

We consider the problems of multiclass (or polytomous) linear classification and feature selection. In both problems, one is given training data of the form $\{(y_m, \mathbf{a}_m)\}_{m=1}^M$, where $\mathbf{a}_m \in \mathbb{R}^N$ is a vector of features and $y_m \in \{1, \dots, D\}$ is the corresponding D -ary class label. In *multiclass classification*, the goal is to infer the unknown label y_0 associated with a newly observed feature vector \mathbf{a}_0 . In the *linear* approach to this problem, the training data are used to design a weight matrix $\mathbf{X} \in \mathbb{R}^{N \times D}$ that generates a vector of “scores” $\mathbf{z}_0 \triangleq \mathbf{X}^T \mathbf{a}_0 \in \mathbb{R}^D$, the largest of which can be used to predict the unknown label, i.e.,

$$\hat{y}_0 = \arg \max_d [\mathbf{z}_0]_d. \quad (1)$$

In *feature selection*, the goal is to determine which *subset* of the N features \mathbf{a}_0 is needed to accurately predict the label y_0 .

We are particularly interested in the setting where the number of features, N , is large and greatly exceeds the number of training examples, M . Such problems arise in a number of important applications, such as micro-array gene expression [1,2], multi-voxel pattern analysis (MVPA) [3,4], text mining [5,6], and analysis of marketing data [7].

The authors are with the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, OH.

Please direct all correspondence to Prof. Philip Schniter, Dept. ECE, 2015 Neil Ave., Columbus OH 43210, e-mail: schniter@ece.osu.edu, phone 614.247.6488, fax 614.292.7596. Evan Byrne can be reached at the same address/phone/fax and e-mailed at byrne.133@osu.edu.

This work was supported in part by the National Science Foundation grants CCF-1018368 and CCF-1218754.

Portions of this work were presented at the 2015 Duke Workshop on Sensing and Analysis of High Dimensional Data.

In the $N \gg M$ case, accurate linear classification and feature selection may be possible if the labels are influenced by a sufficiently small number, K , of the total N features. For example, in binary linear classification, performance guarantees are possible with only $M = O(K \log N/K)$ training examples when \mathbf{a}_m is i.i.d. Gaussian [8]. Note that, when $K \ll N$, accurate linear classification can be accomplished using a *sparse* weight matrix \mathbf{X} , i.e., a matrix where all but a few rows are zero-valued.

B. Multinomial logistic regression

For multiclass linear classification and feature selection, we focus on the approach known as *multinomial logistic regression* (MLR) [9], which can be described using a generative probabilistic model. Here, the label vector $\mathbf{y} \triangleq [y_0, \dots, y_M]^T$ is modeled as a realization of a random¹ vector $\mathbf{Y} \triangleq [Y_0, \dots, Y_M]^T$, the “true” weight matrix \mathbf{X} is modeled as a realization of a random matrix \mathbf{X} , and the features $\mathbf{A} \triangleq [\mathbf{a}_0, \dots, \mathbf{a}_M]^T$ are treated as deterministic. Moreover, the labels y_m are modeled as conditionally independent given the scores $\mathbf{z}_m \triangleq \mathbf{X}^T \mathbf{a}_m$, i.e.,

$$\Pr\{\mathbf{y} = \mathbf{y} | \mathbf{X} = \mathbf{X}; \mathbf{A}\} = \prod_{m=1}^M p_{\mathbf{y}|\mathbf{z}}(y_m | \mathbf{X}^T \mathbf{a}_m), \quad (2)$$

and distributed according to the multinomial logistic (or softmax) pmf:

$$p_{\mathbf{y}|\mathbf{z}}(y_m | \mathbf{z}_m) = \frac{\exp([\mathbf{z}_m]_{y_m})}{\sum_{d=1}^D \exp([\mathbf{z}_m]_d)}, \quad y_m \in \{1, \dots, D\}. \quad (3)$$

The rows \mathbf{x}_n^T of the weight matrix \mathbf{X} are then modeled as i.i.d.,

$$p_{\mathbf{X}}(\mathbf{X}) = \prod_{n=1}^N p_{\mathbf{X}}(\mathbf{x}_n), \quad (4)$$

where $p_{\mathbf{X}}$ may be chosen to promote sparsity.

C. Existing methods

Several sparsity-promoting MLR algorithms have been proposed (e.g., [10,11,12,13,14,15]), differing in their choice of $p_{\mathbf{X}}$ and methodology of estimating \mathbf{X} . For example, [11,12,13] use the i.i.d. Laplacian prior

$$p_{\mathbf{X}}(\mathbf{x}_n; \lambda) = \prod_{d=1}^D \frac{\lambda}{2} \exp(-\lambda |x_{nd}|), \quad (5)$$

¹For clarity, we typeset random quantities in sans-serif font and deterministic quantities in serif font.

with λ tuned via cross-validation. To circumvent this tuning problem, [14] employs the Laplacian scale mixture

$$p_{\mathbf{x}}(\mathbf{x}_n) = \prod_{d=1}^D \int \left[\frac{\lambda}{2} \exp(-\lambda |x_{nd}|) \right] p(\lambda) d\lambda, \quad (6)$$

with Jeffrey's non-informative hyperprior $p(\lambda) \propto \frac{1}{\lambda} 1_{\lambda \geq 0}$. The relevance vector machine (RVM) approach [10] uses the Gaussian scale mixture

$$p_{\mathbf{x}}(\mathbf{x}_n) = \prod_{d=1}^D \int \mathcal{N}(x_{nd}; 0, \nu) p(\nu) d\nu, \quad (7)$$

with inverse-gamma $p(\nu)$ (i.e., the conjugate hyperprior), resulting in an i.i.d. student's t distribution for $p_{\mathbf{x}}$. However, other choices are possible. For example, the exponential hyperprior $p(\nu; \lambda) = \frac{\lambda^2}{2} \exp(-\frac{\lambda^2}{2}\nu) 1_{\nu \geq 0}$ would lead back to the i.i.d. Laplacian distribution (5) for $p_{\mathbf{x}}$ [16]. Finally, [15] uses

$$p_{\mathbf{x}}(\mathbf{x}_n; \lambda) \propto \exp(-\lambda \|\mathbf{x}_n\|_2), \quad (8)$$

which encourages row-sparsity in \mathbf{X} .

Once the probabilistic model (2)-(4) has been specified, a procedure is needed to infer the weights \mathbf{X} from the training data $\{(y_m, \mathbf{a}_m)\}_{m=1}^M$. The Laplacian-prior methods [11,12,13,15] use the maximum a posteriori (MAP) estimation framework:

$$\widehat{\mathbf{X}} = \arg \max_{\mathbf{X}} \log p(\mathbf{X}|\mathbf{y}; \mathbf{A}) \quad (9)$$

$$= \arg \max_{\mathbf{X}} \sum_{m=1}^M \log p_{y|\mathbf{z}}(y_m | \mathbf{X}^T \mathbf{a}_m) + \sum_{n=1}^N \log p_{\mathbf{x}}(\mathbf{x}_n), \quad (10)$$

where Bayes' rule was used for (10). Under $p_{\mathbf{x}}$ from (5) or (8), the second term in (10) reduces to $-\lambda \sum_{n=1}^N \|\mathbf{x}_n\|_1$ or $-\lambda \sum_{n=1}^N \|\mathbf{x}_n\|_2$, respectively. In this case, (10) is concave and can be maximized in polynomial time; [11,12,13,15] employ (block) coordinate ascent for this purpose. The papers [10] and [14] handle the scale-mixture priors (6) and (7), respectively, using the evidence maximization framework [17]. This approach yields a double-loop procedure: the hyperparameter λ or ν is estimated in the outer loop, and—for fixed λ or ν —the resulting concave (i.e., ℓ_2 or ℓ_1 regularized) MAP optimization is solved in the inner loop.

The methods [10,11,12,13,14,15] described above all yield a sparse point estimate $\widehat{\mathbf{X}}$. Thus, feature selection is accomplished by examining the row-support of $\widehat{\mathbf{X}}$ and classification is accomplished through (1).

D. Contributions

In Section II, we propose new approaches to sparse-weight MLR based on the *hybrid generalized approximate message passing* (HyGAMP) framework from [18]. HyGAMP offers tractable approximations of the sum-product and min-sum message passing algorithms [19] by leveraging results of the central limit theorem that hold in the large-system limit: $\lim_{N,M \rightarrow \infty}$ with fixed N/M . Without approximation, both the sum-product algorithm (SPA) and min-sum algorithm (MSA) are intractable due to the forms of $p_{y|\mathbf{z}}$ and $p_{\mathbf{x}}$ in our problem.

For context, we note that HyGAMP is a generalization of the original GAMP approach from [20], which cannot be directly applied to the MLR problem because the likelihood function (3) is not separable, i.e., $p_{y|\mathbf{z}}(y_m | \mathbf{z}_m) \neq \prod_d p(y_m | z_{md})$. GAMP can, however, be applied to *binary* classification and feature selection, as in [21]. Meanwhile, GAMP is itself a generalization of the original AMP approach from [22,23], which requires $p_{y|\mathbf{z}}$ to be both separable and Gaussian.

With the HyGAMP algorithm from [18], message passing for sparse-weight MLR reduces to an iterative update of $O(M+N)$ multivariate Gaussian pdfs, each of dimension D . Although HyGAMP makes MLR tractable, it is still not computationally practical for the large values of M and N in contemporary applications (e.g., $N \sim 10^4$ to 10^6 in genomics and MVPA). Similarly, the non-conjugate variational message passing technique from [24] requires the update of $O(MN)$ multivariate Gaussian pdfs of dimension D , which is even less practical for large M and N .

Thus, in Section III, we propose a simplified HyGAMP (SHyGAMP) algorithm for MLR that approximates HyGAMP's mean and variance computations in an efficient manner. In particular, we investigate approaches based on numerical integration, importance sampling, Taylor-series approximation, and a novel Gaussian-mixture approximation, and we conduct numerical experiments that suggest the superiority of the latter.

In Section IV, we detail two approaches to tune the hyperparameters that control the statistical models assumed by SHyGAMP, one based on the expectation-maximization (EM) methodology from [25] and the other based on a variation of the Stein's unbiased risk estimate (SURE) methodology from [26]. We also give numerical evidence that these methods yield near-optimal hyperparameter estimates.

Finally, in Section V, we compare our proposed SHyGAMP methods to the state-of-the-art MLR approaches [13,14] on both synthetic and practical real-world problems. Our experiments suggest that our proposed methods offer simultaneous improvements in classification error rate and runtime.

Notation: Random quantities are typeset in sans-serif (e.g., \mathbf{x}) while deterministic quantities are typeset in serif (e.g., x). The pdf of random variable \mathbf{x} under deterministic parameters $\boldsymbol{\theta}$ is written as $p_{\mathbf{x}}(x; \boldsymbol{\theta})$, where the subscript and parameterization are sometimes omitted for brevity. Column vectors are typeset in boldface lower-case (e.g., \mathbf{y} or \mathbf{y}), matrices in boldface upper-case (e.g., \mathbf{X} or \mathbf{X}), and their transpose is denoted by $(\cdot)^T$. $E\{\cdot\}$ denotes expectation and $\text{Cov}\{\cdot\}$ autocovariance. \mathbf{I}_K denotes the $K \times K$ identity matrix, \mathbf{e}_k the k th column of \mathbf{I}_K , $\mathbf{1}_K$ the length- K vector of ones, and $\text{Diag}(\mathbf{b})$ the diagonal matrix created from the vector \mathbf{b} . $[\mathbf{B}]_{m,n}$ denotes the element in the m^{th} row and n^{th} column of \mathbf{B} , and $\|\cdot\|_F$ the Frobenius norm. Finally, δ_n denotes the Kronecker delta sequence, $\delta(x)$ the Dirac delta distribution, and 1_A the indicator function of the event A .

II. HYGAMP FOR MULTICLASS CLASSIFICATION

In this section, we detail the application of HyGAMP [18] to multiclass linear classification. In particular, we show that the




Fig. 1: Factor graph representations of (14), with white/gray circles denoting unobserved/observed random variables, and gray rectangles denoting pdf “factors”.

sum-product algorithm (SPA) variant of HyGAMP is a loopy belief propagation (LBP) approximation of the classification-error-rate minimizing linear classifier and that the min-sum algorithm (MSA) variant is an LBP approach to solving the MAP problem (10).

A. Classification via sum-product HyGAMP

Suppose that we are given M labeled training pairs $\{(y_m, \mathbf{a}_m)\}_{m=1}^M$ and T test feature vectors $\{\mathbf{a}_t\}_{t=M+1}^{M+T}$ associated with unknown test labels $\{y_t\}_{t=M+1}^{M+T}$, all obeying the MLR statistical model (2)-(4). Consider the problem of computing the classification-error-rate minimizing hypotheses $\{\hat{y}_t\}_{t=M+1}^{M+T}$,

$$\hat{y}_t = \arg \max_{y_t \in \{1, \dots, D\}} p_{y_t | \mathbf{y}_{1:M}}(y_t | \mathbf{y}_{1:M}; \mathbf{A}), \quad (11)$$

under known $p_{y|\mathbf{z}}$ and $p_{\mathbf{x}}$, where $\mathbf{y}_{1:M} \triangleq [y_1, \dots, y_M]^T$ and $\mathbf{A} \triangleq [\mathbf{a}_1, \dots, \mathbf{a}_{M+T}]^T$. The probabilities in (11) can be computed via the marginalization

$$p_{y_t | \mathbf{y}_{1:M}}(y_t | \mathbf{y}_{1:M}; \mathbf{A}) = p_{y_t, \mathbf{y}_{1:M}}(y_t, \mathbf{y}_{1:M}; \mathbf{A}) Z_{\mathbf{y}}^{-1} \quad (12)$$

$$= Z_{\mathbf{y}}^{-1} \sum_{\mathbf{y} \in \mathcal{Y}_t(y_t)} \int p_{\mathbf{y}, \mathbf{x}}(\mathbf{y}, \mathbf{x}; \mathbf{A}) d\mathbf{x}, \quad (13)$$

with scaling constant $Z_{\mathbf{y}}^{-1}$, label vector $\mathbf{y} = [y_1, \dots, y_{M+T}]^T$, and constraint set $\mathcal{Y}_t(y_t) \triangleq \{\tilde{\mathbf{y}} \in \{1, \dots, D\}^{M+T} \text{ s.t. } [\tilde{\mathbf{y}}]_t = y_t \text{ and } [\tilde{\mathbf{y}}]_m = y_m \forall m = 1, \dots, M\}$, which fixes the t th element of \mathbf{y} at the value y_t and the first M elements of \mathbf{y} at the values of the corresponding training labels. Due to (2) and (4), the joint pdf in (13) factors as

$$p_{\mathbf{y}, \mathbf{x}}(\mathbf{y}, \mathbf{x}; \mathbf{A}) = \prod_{m=1}^{M+T} p_{y|\mathbf{z}}(y_m | \mathbf{X}^T \mathbf{a}_m) \prod_{n=1}^N p_{\mathbf{x}}(\mathbf{x}_n). \quad (14)$$

The factorization in (14) is depicted by the *factor graph* in Fig. 1a, where the random variables $\{y_m\}$ and random vectors $\{\mathbf{x}_n\}$ are connected to the pdf factors in which they appear.

Since exact computation of the marginal posterior test-label probabilities is an NP-hard problem [27], we are interested in alternative strategies, such as those based on loopy belief propagation by the SPA [19]. Although a direct application of the SPA is itself intractable when $p_{y|\mathbf{z}}$ takes the MLR form (3), the SPA simplifies in the large-system limit under i.i.d. sub-Gaussian \mathbf{A} , leading to the HyGAMP approximation [18]

given² in Algorithm 1. Although in practical MLR applications \mathbf{A} is not i.i.d. Gaussian,³ the numerical results in Section V suggest that treating it as such works sufficiently well.

We note from Fig. 1a that the HyGAMP algorithm is applicable to a factor graph with vector-valued variable nodes. As such, it generalizes the GAMP algorithm from [20], which applies only to a factor graph with scalar-variable nodes. Below, we give a brief explanation for the steps in Algorithm 1. For those interested in more details, we suggest [18] for an overview and derivation of HyGAMP, [20] for an overview and derivation of GAMP, [28] for rigorous analysis of GAMP under large i.i.d. sub-Gaussian \mathbf{A} , and [29,30] for fixed-point and local-convergence analysis of GAMP under arbitrary \mathbf{A} .

Lines 6-7 of Algorithm 1 produce an approximation of the posterior mean and covariance of \mathbf{x}_n at each iteration t . Similarly, lines 15-16 produce an approximation of the posterior mean and covariance of $\mathbf{z}_m \triangleq \mathbf{X}^T \mathbf{a}_m$. The posterior mean and covariance of \mathbf{x}_n are computed from the intermediate quantity $\hat{\mathbf{r}}_n(t)$, which behaves like a noisy measurement of the true \mathbf{x}_n . In particular, for i.i.d. Gaussian \mathbf{A} in the large-system limit, $\hat{\mathbf{r}}_n(t)$ is a typical realization of the random vector $\mathbf{r}_n = \mathbf{x}_n + \mathbf{v}_n$ with $\mathbf{v}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_n^{\mathbf{r}}(t))$. Thus, the approximate posterior pdf used in lines 6-7 is

$$p_{\mathbf{x}|\mathbf{r}}(\mathbf{x}_n | \hat{\mathbf{r}}_n; \mathbf{Q}_n^{\mathbf{r}}) = \frac{p_{\mathbf{x}}(\mathbf{x}_n) \mathcal{N}(\mathbf{x}_n; \hat{\mathbf{r}}_n, \mathbf{Q}_n^{\mathbf{r}})}{\int p_{\mathbf{x}}(\mathbf{x}'_n) \mathcal{N}(\mathbf{x}'_n; \hat{\mathbf{r}}_n, \mathbf{Q}_n^{\mathbf{r}}) d\mathbf{x}'_n}. \quad (15)$$

A similar interpretation holds for HyGAMP's approximation of the posterior mean and covariance of \mathbf{z}_m in lines 15-16, which uses the intermediate vector $\hat{\mathbf{p}}_m(t)$ and the approximate posterior pdf

$$p_{\mathbf{z}|\mathbf{y}, \mathbf{p}}(\mathbf{z}_m | y_m, \hat{\mathbf{p}}_m; \mathbf{Q}_m^{\mathbf{p}}) = \frac{p_{y|\mathbf{z}}(y_m | \mathbf{z}_m) \mathcal{N}(\mathbf{z}_m; \hat{\mathbf{p}}_m, \mathbf{Q}_m^{\mathbf{p}})}{\int p_{y|\mathbf{z}}(y_m | \mathbf{z}'_m) \mathcal{N}(\mathbf{z}'_m; \hat{\mathbf{p}}_m, \mathbf{Q}_m^{\mathbf{p}}) d\mathbf{z}'_m}. \quad (16)$$

B. Classification via min-sum HyGAMP

As discussed in Section I-C, an alternative approach to linear classification and feature selection is through MAP estimation of the true weight matrix \mathbf{X} . Given a likelihood of the form (2) and a prior of the form (4), the MAP estimate is the solution to the optimization problem (10).

Similar to how the SPA can be used to compute approximate marginal posteriors in loopy graphs, the min-sum algorithm (MSA) [19] can be used to compute the MAP estimate. Although a direct application of the MSA is intractable when $p_{y|\mathbf{z}}$ takes the MLR form (3), the MSA simplifies in the large-system limit under i.i.d. sub-Gaussian \mathbf{A} , leading to the MSA form of HyGAMP specified in Algorithm 1.

As described in Section II-A, when \mathbf{A} is large and i.i.d. sub-Gaussian, the vector $\hat{\mathbf{r}}_n(t)$ in Algorithm 1 behaves like a Gaussian-noise-corrupted observation of the true \mathbf{x}_n with noise covariance $\mathbf{Q}_n^{\mathbf{r}}(t)$. Thus, line 3 can be interpreted as

²The HyGAMP algorithm in [18] is actually more general than what is specified in Algorithm 1, but the version in Algorithm 1 is sufficient to handle the factor graph in Fig. 1a.

³We note that many of the standard data pre-processing techniques, such as z-scoring, tend to make the feature distributions closer to zero-mean Gaussian.

Algorithm 1 HyGAMP

Require: Mode $\in \{\text{SPA}, \text{MSA}\}$, matrix \mathbf{A} , vector \mathbf{y} , pdfs $p_{\mathbf{x}|\mathbf{r}}$ and $p_{\mathbf{z}|\mathbf{y}, \mathbf{p}}$ from (15)-(16), initializations $\hat{\mathbf{r}}_n(0)$, $\mathbf{Q}_n^{\mathbf{r}}(0)$.
Ensure: $t \leftarrow 0$; $\hat{\mathbf{s}}_m(0) \leftarrow \mathbf{0}$.
 1: **repeat**
 2: **if** MSA **then** **for** $n = 1 \dots N$
 3: $\hat{\mathbf{x}}_n(t) \leftarrow \arg \max_{\mathbf{x}} \log p_{\mathbf{x}|\mathbf{r}}(\mathbf{x}_n | \hat{\mathbf{r}}_n(t-1); \mathbf{Q}_n^{\mathbf{r}}(t-1))$
 4: $\mathbf{Q}_n^{\mathbf{x}}(t) \leftarrow [-\frac{\partial^2}{\partial \mathbf{x}^2} \log p_{\mathbf{x}|\mathbf{r}}(\hat{\mathbf{x}}_n(t) | \hat{\mathbf{r}}_n(t-1); \mathbf{Q}_n^{\mathbf{r}}(t-1))]^{-1}$
 5: **else if** SPA **then** **for** $n = 1 \dots N$
 6: $\hat{\mathbf{x}}_n(t) \leftarrow \mathbb{E} \{\mathbf{x}_n | \mathbf{r}_n = \hat{\mathbf{r}}_n(t-1); \mathbf{Q}_n^{\mathbf{r}}(t-1)\}$
 7: $\mathbf{Q}_n^{\mathbf{x}}(t) \leftarrow \text{Cov} \{\mathbf{x}_n | \mathbf{r}_n = \hat{\mathbf{r}}_n(t-1); \mathbf{Q}_n^{\mathbf{r}}(t-1)\}$
 8: **end if**
 9: $\forall m : \mathbf{Q}_m^{\mathbf{p}}(t) \leftarrow \sum_{n=1}^N A_{mn}^2 \mathbf{Q}_n^{\mathbf{x}}(t)$
 10: $\forall m : \hat{\mathbf{p}}_m(t) \leftarrow \sum_{n=1}^N A_{mn} \hat{\mathbf{x}}_n(t) - \mathbf{Q}_m^{\mathbf{p}}(t) \hat{\mathbf{s}}_m(t-1)$
 11: **if** MSA **then** **for** $m = 1 \dots M$
 12: $\hat{\mathbf{z}}_m(t) \leftarrow \arg \max_{\mathbf{z}} \log p_{\mathbf{z}|\mathbf{y}, \mathbf{p}}(\mathbf{z}_m | y_m, \hat{\mathbf{p}}_m(t); \mathbf{Q}_m^{\mathbf{p}}(t))$
 13: $\mathbf{Q}_m^{\mathbf{z}}(t) \leftarrow [-\frac{\partial^2}{\partial \mathbf{z}^2} \log p_{\mathbf{z}|\mathbf{y}, \mathbf{p}}(\hat{\mathbf{z}}_m(t) | y_m, \hat{\mathbf{p}}_m(t); \mathbf{Q}_m^{\mathbf{p}}(t))]^{-1}$
 14: **else if** SPA **then** **for** $m = 1 \dots M$
 15: $\hat{\mathbf{z}}_m(t) \leftarrow \mathbb{E} \{\mathbf{z}_m | y_m, \mathbf{p}_m = \hat{\mathbf{p}}_m(t); \mathbf{Q}_m^{\mathbf{p}}(t)\}$
 16: $\mathbf{Q}_m^{\mathbf{z}}(t) \leftarrow \text{Cov} \{\mathbf{z}_m | y_m, \mathbf{p}_m = \hat{\mathbf{p}}_m(t); \mathbf{Q}_m^{\mathbf{p}}(t)\}$
 17: **end if**
 18: $\forall m : \mathbf{Q}_m^{\mathbf{s}}(t) \leftarrow [\mathbf{Q}_m^{\mathbf{p}}(t)]^{-1} - [\mathbf{Q}_m^{\mathbf{p}}(t)]^{-1} \mathbf{Q}_m^{\mathbf{z}}(t) [\mathbf{Q}_m^{\mathbf{p}}(t)]^{-1}$
 19: $\forall m : \hat{\mathbf{s}}_m(t) \leftarrow [\mathbf{Q}_m^{\mathbf{p}}(t)]^{-1} (\hat{\mathbf{z}}_m(t) - \hat{\mathbf{p}}_m(t))$
 20: $\forall n : \mathbf{Q}_n^{\mathbf{r}}(t) \leftarrow [\sum_{m=1}^M A_{mn}^2 \mathbf{Q}_m^{\mathbf{s}}(t)]^{-1}$
 21: $\forall n : \hat{\mathbf{r}}_n(t) \leftarrow \hat{\mathbf{x}}_n(t) + \mathbf{Q}_n^{\mathbf{r}}(t) \sum_{m=1}^M A_{mn} \hat{\mathbf{s}}_m(t)$
 22: $t \leftarrow t + 1$
 23: **until** Terminated

MAP estimation of \mathbf{x}_n and line 4 as measuring the local curvature of the corresponding MAP cost. Similar interpretations hold for MAP estimation of \mathbf{z}_m via lines 12-13.

C. Implementation of sum-product HyGAMP

From Algorithm 1, we see that HyGAMP requires inverting $M+N$ matrices of size $D \times D$ (for lines 18 and 20) in addition to solving $M+N$ joint inference problems of dimension D in lines 3-7 and 12-16. We now briefly discuss the latter problems for the sum-product version of HyGAMP.

1) *Inference of \mathbf{x}_n* : One choice of weight-coefficient prior $p_{\mathbf{x}_n}$ that facilitates row-sparse \mathbf{X} and tractable SPA inference is Bernoulli-multivariate-Gaussian, i.e.,

$$p_{\mathbf{x}}(\mathbf{x}_n) = (1 - \beta) \delta(\mathbf{x}_n) + \beta \mathcal{N}(\mathbf{x}_n; \mathbf{0}, v\mathbf{I}), \quad (17)$$

where $\delta(\cdot)$ denotes the Dirac delta and $\beta \in (0, 1]$. In this case, it can be shown [31] that the mean and variance computations in lines 6-7 of Algorithm 1 reduce to

$$C_n = 1 + \frac{1 - \beta}{\beta} \frac{\mathcal{N}(\mathbf{0}; \hat{\mathbf{r}}_n, \mathbf{Q}_n^{\mathbf{r}})}{\mathcal{N}(\mathbf{0}; \hat{\mathbf{r}}_n, v\mathbf{I} + \mathbf{Q}_n^{\mathbf{r}})} \quad (18)$$

$$\hat{\mathbf{x}}_n = C_n^{-1} (\mathbf{I} + v^{-1} \mathbf{Q}_n^{\mathbf{r}})^{-1} \hat{\mathbf{r}}_n \quad (19)$$

$$\mathbf{Q}_n^{\mathbf{x}} = C_n^{-1} (\mathbf{I} + v^{-1} \mathbf{Q}_n^{\mathbf{r}})^{-1} \mathbf{Q}_n^{\mathbf{r}} + (C_n - 1) \hat{\mathbf{x}}_n \hat{\mathbf{x}}_n^{\mathbf{T}}, \quad (20)$$

which requires a $D \times D$ matrix inversion at each n .

2) *Inference of \mathbf{z}_m* : When $p_{\mathbf{y}|\mathbf{z}}$ takes the MLR form in (3), closed-form expressions for $\hat{\mathbf{z}}_m(t)$ and $\mathbf{Q}_m^{\mathbf{z}}(t)$ from lines 15-16 of Algorithm 1 do not exist. While these computations could be approximated using, e.g., numerical integration or importance sampling, this is expensive because $\hat{\mathbf{z}}_m(t)$ and $\mathbf{Q}_m^{\mathbf{z}}(t)$ must be computed for every index m at every HyGAMP

iteration t . More details on these approaches will be presented in Section III-C, in the context of SHyGAMP.

D. Implementation of min-sum HyGAMP

1) *Inference of \mathbf{x}_n* : To ease the computation of line 3 in Algorithm 1, it is typical to choose a log-concave prior $p_{\mathbf{x}}$ so that the optimization problem (10) is concave (since $p_{\mathbf{y}|\mathbf{z}}$ in (3) is also log-concave). As discussed in Section I-C, a common example of a log-concave sparsity-promoting prior is the Laplace prior (5). In this case, line 3 becomes

$$\hat{\mathbf{x}}_n = \arg \max_{\mathbf{x}} -\frac{1}{2} (\mathbf{x} - \hat{\mathbf{r}}_n)^{\mathbf{T}} [\mathbf{Q}_n^{\mathbf{r}}]^{-1} (\mathbf{x} - \hat{\mathbf{r}}_n) - \lambda \|\mathbf{x}\|_1, \quad (21)$$

which is essentially the LASSO [32] problem. Although (21) has no closed-form solution, it can be solved iteratively using, e.g., minorization-maximization (MM) [33].

To maximize a function $J(\mathbf{x})$, MM iterates the recursion

$$\hat{\mathbf{x}}^{(k+1)} = \arg \max_{\mathbf{x}} \hat{J}(\mathbf{x}; \hat{\mathbf{x}}^{(k)}), \quad (22)$$

where $\hat{J}(\mathbf{x}; \hat{\mathbf{x}})$ is a surrogate function that minorizes $J(\mathbf{x})$ at $\hat{\mathbf{x}}$. In other words, $\hat{J}(\mathbf{x}; \hat{\mathbf{x}}) \leq J(\mathbf{x}) \forall \mathbf{x}$ for any fixed $\hat{\mathbf{x}}$, with equality when $\mathbf{x} = \hat{\mathbf{x}}$. To apply MM to (21), we identify the utility function as $J_n(\mathbf{x}) \triangleq -\frac{1}{2} (\mathbf{x} - \hat{\mathbf{r}}_n)^{\mathbf{T}} [\mathbf{Q}_n^{\mathbf{r}}]^{-1} (\mathbf{x} - \hat{\mathbf{r}}_n) - \lambda \|\mathbf{x}\|_1$. Next we apply a result from [34] that established that $J_n(\mathbf{x})$ is minorized by $\hat{J}_n(\mathbf{x}; \hat{\mathbf{x}}_n^{(k)}) \triangleq -\frac{1}{2} (\mathbf{x} - \hat{\mathbf{r}}_n)^{\mathbf{T}} [\mathbf{Q}_n^{\mathbf{r}}]^{-1} (\mathbf{x} - \hat{\mathbf{r}}_n) - \frac{\lambda}{2} (\mathbf{x}^{\mathbf{T}} \mathbf{\Lambda}(\hat{\mathbf{x}}_n^{(k)}) \mathbf{x} + \|\hat{\mathbf{x}}_n^{(k)}\|_2^2)$ with $\mathbf{\Lambda}(\hat{\mathbf{x}}) \triangleq \text{Diag} \{|\hat{x}_1|^{-1}, \dots, |\hat{x}_D|^{-1}\}$. Thus (22) implies

$$\hat{\mathbf{x}}_n^{(k+1)} = \arg \max_{\mathbf{x}} \hat{J}_n(\mathbf{x}; \hat{\mathbf{x}}_n^{(k)}) \quad (23)$$

$$= \arg \max_{\mathbf{x}} \mathbf{x}^{\mathbf{T}} [\mathbf{Q}_n^{\mathbf{r}}]^{-1} \hat{\mathbf{r}}_n - \frac{1}{2} \mathbf{x}^{\mathbf{T}} ([\mathbf{Q}_n^{\mathbf{r}}]^{-1} + \lambda \mathbf{\Lambda}(\hat{\mathbf{x}}_n^{(k)})) \mathbf{x} \quad (24)$$

$$= ([\mathbf{Q}_n^{\mathbf{r}}]^{-1} + \lambda \mathbf{\Lambda}(\hat{\mathbf{x}}_n^{(k)}))^{-1} [\mathbf{Q}_n^{\mathbf{r}}]^{-1} \hat{\mathbf{r}}_n \quad (25)$$

where (24) dropped the \mathbf{x} -invariant terms from $\hat{J}_n(\mathbf{x}; \hat{\mathbf{x}}_n^{(k)})$. Note that each iteration k of (25) requires a $D \times D$ matrix inverse for each n .

Line 4 of Algorithm 1 then says to set $\mathbf{Q}_n^{\mathbf{x}}$ equal to the Hessian of the objective function in (21) at $\hat{\mathbf{x}}_n$. Recalling that the second derivative of $|\mathbf{x}_{nd}|$ is undefined when $\mathbf{x}_{nd} = 0$ but otherwise equals zero, we set $\mathbf{Q}_n^{\mathbf{x}} = \mathbf{Q}_n^{\mathbf{r}}$ but then zero the d th row and column of $\mathbf{Q}_n^{\mathbf{x}}$ for all d such that $\hat{\mathbf{x}}_{nd} = 0$.

2) *Inference of \mathbf{z}_m* : Min-sum HyGAMP also requires the computation of lines 12-13 in Algorithm 1. In our MLR application, line 12 reduces to the concave optimization problem

$$\hat{\mathbf{z}}_m = \arg \max_{\mathbf{z}} -\frac{1}{2} (\mathbf{z} - \hat{\mathbf{p}}_m)^{\mathbf{T}} [\mathbf{Q}_m^{\mathbf{p}}]^{-1} (\mathbf{z} - \hat{\mathbf{p}}_m) + \log p_{\mathbf{y}|\mathbf{z}}(y_m | \mathbf{z}). \quad (26)$$

Although (26) can be solved in a variety of ways (see [31] for MM-based methods), we now describe one based on Newton's method [35], i.e.,

$$\hat{\mathbf{z}}_m^{(k+1)} = \hat{\mathbf{z}}_m^{(k)} - \alpha^{(k)} [\mathbf{H}_m^{(k)}]^{-1} \mathbf{g}_m^{(k)}, \quad (27)$$

where $\mathbf{g}_m^{(k)}$ and $\mathbf{H}_m^{(k)}$ are the gradient and Hessian of the objective function in (26) at $\hat{\mathbf{z}}_m^{(k)}$, and $\alpha^{(k)} \in (0, 1]$ is a

stepsize. From (3), it can be seen that $\frac{\partial}{\partial z_i} \log p_{y|z}(y|z) = \delta_{y-i} - p_{y|z}(i|z)$, and so

$$\mathbf{g}_m^{(k)} = \mathbf{u}(\hat{\mathbf{z}}_m^{(k)}) - \mathbf{e}_{y_m} + [\mathbf{Q}_m^{\mathbf{p}}]^{-1}(\hat{\mathbf{z}}_m^{(k)} - \hat{\mathbf{p}}_m), \quad (28)$$

where \mathbf{e}_y denotes the y th column of \mathbf{I}_D and $\mathbf{u}(z) \in \mathbb{R}^{D \times 1}$ is defined elementwise as

$$[\mathbf{u}(z)]_i \triangleq p_{y|z}(i|z). \quad (29)$$

Similarly, it is known [36] that the Hessian takes the form

$$\mathbf{H}_m^{(k)} = \mathbf{u}(\hat{\mathbf{z}}_m) \mathbf{u}(\hat{\mathbf{z}}_m)^{\top} - \text{Diag}\{\mathbf{u}(\hat{\mathbf{z}}_m)\} - [\mathbf{Q}_m^{\mathbf{p}}]^{-1}, \quad (30)$$

which also provides the answer to line 13 of Algorithm 1. Note that each iteration k of (27) requires a $D \times D$ matrix inverse for each m .

It is possible to circumvent the matrix inversion in (27) via componentwise update, i.e.,

$$\hat{z}_{md}^{(k+1)} = \hat{z}_{md}^{(k)} - \alpha^{(k)} g_{md}^{(k)} / H_{md}^{(k)}, \quad (31)$$

where $g_{md}^{(k)}$ and $H_{md}^{(k)}$ are the first and second derivatives of the objective function in (26) with respect to z_d at $\mathbf{z} = \hat{\mathbf{z}}_m^{(k)}$. From (28)-(30), it follows that

$$g_{md}^{(k)} = p_{y|z}(d|\hat{z}_m^{(k)}) - \delta_{y_m-d} + [[\mathbf{Q}_m^{\mathbf{p}}]^{-1}]_{:,d}^{\top} (\hat{\mathbf{z}}_m^{(k)} - \hat{\mathbf{p}}_m) \quad (32)$$

$$H_{md}^{(k)} = p_{y|z}(d|\hat{z}_m^{(k)})^2 - p_{y|z}(d|\hat{z}_m^{(k)}) - [[\mathbf{Q}_m^{\mathbf{p}}]^{-1}]_{dd}. \quad (33)$$

E. HyGAMP summary

In summary, the SPA and MSA variants of the HyGAMP algorithm provide tractable methods of approximating the posterior test-label probabilities $p_{y_t|y_{1:M}}(y_t | \mathbf{y}_{1:M}; \mathbf{A})$ and computing the MAP weight matrix $\hat{\mathbf{X}} = \arg \max_{\mathbf{X}} p_{\mathbf{y}_{1:M}, \mathbf{X}}(\mathbf{y}_{1:M}; \mathbf{X}; \mathbf{A})$, respectively, under a separable likelihood (2) and a separable prior (4). In particular, HyGAMP attacks the high-dimensional inference problems of interest using a sequence of $M + N$ low-dimensional (in particular, D -dimensional) inference problems and $D \times D$ matrix inversions, as detailed in Algorithm 1.

As detailed in the previous subsections, however, these D -dimensional inference problems are non-trivial in the sparse MLR case, making HyGAMP computationally costly. We refer the reader to Table I for a summary of the D -dimensional inference problems encountered in running SPA-HyGAMP or MSA-HyGAMP, as well as their associated computational costs. Thus, in the sequel, we propose a computationally efficient simplification of HyGAMP that, as we will see in Section V, compares favorably with existing state-of-the-art methods.

III. SHyGAMP FOR MULTICLASS CLASSIFICATION

As described in Section II, a direct application of HyGAMP to sparse MLR is computationally costly. Thus, in this section, we propose a *simplified HyGAMP* (SHyGAMP) algorithm for sparse MLR, whose complexity is greatly reduced. The simplification itself is rather straightforward: we constrain the

Algorithm	Quantity	Method	Complexity
SPA-HyGAMP	$\hat{\mathbf{x}}$	CF	$O(D^3)$
	$\mathbf{Q}^{\mathbf{x}}$	CF	$O(D^3)$
	$\hat{\mathbf{z}}$	NI	$O(D^K)$
	$\mathbf{Q}^{\mathbf{z}}$	NI	$O(D^K)$
MSA-HyGAMP	$\hat{\mathbf{x}}$	MM	$O(KD^3)$
	$\mathbf{Q}^{\mathbf{x}}$	CF	$O(D^3)$
	$\hat{\mathbf{z}}$	CWN	$O(KD^2 + D^3)$
	$\mathbf{Q}^{\mathbf{z}}$	CF	$O(D^3)$

TABLE I: A summary of the D -dimensional inference sub-problems encountered when running SPA-HyGAMP or MSA-HyGAMP, as well as their associated computational costs. ‘CF’ = ‘closed form’, ‘NI’ = ‘numerical integration’, ‘MM’ = ‘minorization-maximization’, and ‘CWN’ = ‘component-wise Newton’s method’. For the NI method, K denotes the number of samples per dimension, and for the MM and CWN methods K denotes the number of iterations.

covariance matrices $\mathbf{Q}_n^{\mathbf{r}}$, $\mathbf{Q}_n^{\mathbf{x}}$, $\mathbf{Q}_m^{\mathbf{p}}$, and $\mathbf{Q}_m^{\mathbf{z}}$ to be diagonal. In other words,

$$\mathbf{Q}_n^{\mathbf{r}} = \text{Diag}\{q_{n1}^{\mathbf{r}}, \dots, q_{nD}^{\mathbf{r}}\}, \quad (34)$$

and similar for $\mathbf{Q}_n^{\mathbf{x}}$, $\mathbf{Q}_m^{\mathbf{p}}$, and $\mathbf{Q}_m^{\mathbf{z}}$. As a consequence, the $D \times D$ matrix inversions in lines 18 and 20 of Algorithm 1 each reduce to D scalar inversions. More importantly, the D -dimensional inference problems in lines 3-7 and 12-16 can be tackled using much simpler methods than those described in Section II, as we detail below.

A. Scalar Variance Approximation

We further approximate the SHyGAMP algorithm using the *scalar variance* GAMP approximation from [18], which reduces the memory and complexity of the algorithm. The scalar variance approximation first approximates the variances $\{q_{nd}^{\mathbf{x}}\}$ by a value invariant to both n and d , i.e.,

$$q^{\mathbf{x}} \triangleq \frac{1}{ND} \sum_{n=1}^N \sum_{d=1}^D q_{nd}^{\mathbf{x}}. \quad (35)$$

Then, in line 9 in Algorithm 1, we use the approximation

$$q_{md}^{\mathbf{p}} \approx \sum_{n=1}^N A_{mn}^2 q^{\mathbf{x}} \stackrel{(a)}{\approx} \frac{\|\mathbf{A}\|_F^2}{M} q^{\mathbf{x}} \triangleq q^{\mathbf{p}}. \quad (36)$$

The approximation (a), after precomputing $\|\mathbf{A}\|_F^2$, reduces the complexity of line 9 from $O(ND)$ to $O(1)$. We next define

$$q^{\mathbf{s}} \triangleq \frac{1}{MD} \sum_{m=1}^M \sum_{d=1}^D q_{md}^{\mathbf{s}} \quad (37)$$

and in line 20 we use the approximation

$$q_{nd}^{\mathbf{r}} \approx \left(\sum_{m=1}^M A_{mn}^2 q^{\mathbf{s}} \right)^{-1} \approx \frac{N}{q^{\mathbf{s}} \|\mathbf{A}\|_F^2} \triangleq q^{\mathbf{r}}. \quad (38)$$

The complexity of line 20 then simplifies from $O(MD)$ to $O(1)$. For clarity, we note that after applying the scalar variance approximation, we have $\mathbf{Q}_n^{\mathbf{x}} = q^{\mathbf{x}} \mathbf{I}_D \forall n$, and similar for $\mathbf{Q}_n^{\mathbf{r}}$, $\mathbf{Q}_m^{\mathbf{p}}$ and $\mathbf{Q}_m^{\mathbf{z}}$.

B. Sum-product SHyGAMP: Inference of x_n

With diagonal \mathbf{Q}_n^r and \mathbf{Q}_n^x , the implementation of lines 6-7 is greatly simplified by choosing a sparsifying prior $p_{\mathbf{x}}$ with the separable form $p_{\mathbf{x}}(\mathbf{x}_n) = \prod_{d=1}^D p_x(x_{nd})$. A common example is the Bernoulli-Gaussian (BG) prior

$$p_x(x_{nd}) = (1 - \beta_d)\delta(x_{nd}) + \beta_d\mathcal{N}(x_{nd}; m_d, v_d\mathbf{I}). \quad (39)$$

For any separable $p_{\mathbf{x}}$, lines 6-7 reduce to computing the mean and variance of the distribution

$$p_{\mathbf{x}|\mathbf{r}}(x_{nd}|\hat{\mathbf{r}}_{nd}; q_{nd}^r) = \frac{p_x(x_{nd})\mathcal{N}(x_{nd}; \hat{\mathbf{r}}_{nd}, q_{nd}^r)}{\int p_x(x'_{nd})\mathcal{N}(x'_{nd}; \hat{\mathbf{r}}_{nd}, q_{nd}^r) dx'_{nd}}. \quad (40)$$

for all $n = 1 \dots N$ and $d = 1 \dots D$, as in the simpler GAMP algorithm [20]. With the BG prior (39), these quantities can be computed in closed form (see, e.g., [37]).

C. Sum-product SHyGAMP: Inference of \mathbf{z}_m

With diagonal \mathbf{Q}_m^p and \mathbf{Q}_m^z , the implementation of lines 15-16 can also be greatly simplified. Essentially, the problem becomes that of computing the scalar means and variances

$$\hat{z}_{md} = C_m^{-1} \int_{\mathbb{R}^D} z_d p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z}) \prod_{k=1}^D \mathcal{N}(z_k; \hat{p}_{mk}, q_{mk}^p) d\mathbf{z} \quad (41)$$

$$q_{md}^z = C_m^{-1} \int_{\mathbb{R}^D} z_d^2 p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z}) \prod_{k=1}^D \mathcal{N}(z_k; \hat{p}_{mk}, q_{mk}^p) d\mathbf{z} - \hat{z}_{md}^2 \quad (42)$$

for $m = 1 \dots M$ and $d = 1 \dots D$. Here, $p_{\mathbf{y}|\mathbf{z}}$ has the MLR form in (3) and C_m is a normalizing constant defined as

$$C_m \triangleq \int_{\mathbb{R}^D} p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z}) \prod_{k=1}^D \mathcal{N}(z_k; \hat{p}_{mk}, q_{mk}^p) d\mathbf{z}. \quad (43)$$

Note that the likelihood $p_{\mathbf{y}|\mathbf{z}}$ is not separable and so inference does not decouple across d , as it did in (40). We now describe several approaches to computing (41)-(42).

1) *Numerical integration*: A straightforward approach to (approximately) computing (41)-(43) is through numerical integration (NI). For this, we propose to use a hyper-rectangular grid of \mathbf{z} values where, for z_d , the interval $[\hat{p}_{md} - \alpha\sqrt{q_{md}^p}, \hat{p}_{md} + \alpha\sqrt{q_{md}^p}]$ is sampled at K equispaced points. Because a D -dimensional numerical integral must be computed for each index m and d , the complexity of this approach grows as $O(MDK^D)$, making it impractical unless D , the number of classes, is very small.

2) *Importance sampling*: An alternative approximation of (41)-(43) can be obtained through importance sampling (IS) [9, §11.1.4]. Here, we draw K independent samples $\{\tilde{\mathbf{z}}_m[k]\}_{k=1}^K$

from $\mathcal{N}(\hat{\mathbf{p}}_m, \mathbf{Q}_m^p)$ and compute

$$C_m \approx \sum_{k=1}^K p_{\mathbf{y}|\mathbf{z}}(y_m|\tilde{\mathbf{z}}_m[k]) \quad (44)$$

$$\hat{z}_{md} \approx C_m^{-1} \sum_{k=1}^K \tilde{z}_{md}[k] p_{\mathbf{y}|\mathbf{z}}(y_m|\tilde{\mathbf{z}}_m[k]) \quad (45)$$

$$q_{md}^z \approx C_m^{-1} \sum_{k=1}^K \tilde{z}_{md}^2[k] p_{\mathbf{y}|\mathbf{z}}(y_m|\tilde{\mathbf{z}}_m[k]) - \hat{z}_{md}^2 \quad (46)$$

for all m and d . The complexity of this approach grows as $O(MDK)$.

3) *Taylor-series approximation*: Another approach is to approximate the likelihood $p_{\mathbf{y}|\mathbf{z}}$ using a second-order Taylor series (TS) about $\hat{\mathbf{p}}_m$, i.e., $p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z}) \approx f_m(\mathbf{z}; \hat{\mathbf{p}}_m)$ with

$$f_m(\mathbf{z}; \hat{\mathbf{p}}_m) \triangleq p_{\mathbf{y}|\mathbf{z}}(y_m|\hat{\mathbf{p}}_m) + \mathbf{g}_m(\hat{\mathbf{p}}_m)^\top (\mathbf{z} - \hat{\mathbf{p}}_m) + \frac{1}{2} (\mathbf{z} - \hat{\mathbf{p}}_m)^\top \mathbf{H}_m(\hat{\mathbf{p}}_m) (\mathbf{z} - \hat{\mathbf{p}}_m) \quad (47)$$

for gradient $\mathbf{g}_m(\hat{\mathbf{p}}) \triangleq \frac{\partial}{\partial \mathbf{z}} p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z})|_{\mathbf{z}=\hat{\mathbf{p}}}$ and Hessian $\mathbf{H}_m(\hat{\mathbf{p}}) \triangleq \frac{\partial^2}{\partial \mathbf{z}^2} p_{\mathbf{y}|\mathbf{z}}(y_m|\mathbf{z})|_{\mathbf{z}=\hat{\mathbf{p}}}$. In this case, it can be shown [31] that

$$C_m \approx f_m(\hat{\mathbf{p}}_m) + \frac{1}{2} \sum_{k=1}^D H_{mk}(\hat{\mathbf{p}}_m) q_{mk}^p \quad (48)$$

$$\hat{z}_{md} \approx \hat{C}_m^{-1} \left(f_m(\hat{\mathbf{p}}_m) \hat{p}_{md} + g_{md}(\hat{\mathbf{p}}_m) q_{md}^p + \frac{1}{2} \sum_{k=1}^D \hat{p}_{mk} q_{mk}^p H_{mk}(\hat{\mathbf{p}}_m) \right) \quad (49)$$

$$q_{md}^z \approx C_m^{-1} \left(f_m(\hat{\mathbf{p}}_m) (\hat{p}_{md}^2 + q_{md}^p) + 2g_{md}(\hat{\mathbf{p}}_m) \hat{p}_{md} q_{md}^p + \frac{1}{2} q_{md}^p (\hat{p}_{md}^2 + 3q_{md}^p) H_{md}(\hat{\mathbf{p}}_m) + \frac{1}{2} (\hat{p}_{md}^2 + q_{md}^p) H_{md}(\hat{\mathbf{p}}_m) \sum_{k \neq d} q_{mk}^p \right) - \hat{z}_{md}^2, \quad (50)$$

where $H_{md}(\hat{\mathbf{p}}) \triangleq [\mathbf{H}_m(\hat{\mathbf{p}})]_{dd}$. The complexity of this approach grows as $O(MD)$.

4) *Gaussian mixture approximation*: It is known that the logistic cdf $1/(1 + \exp(-x))$ is well approximated by a mixture of a few Gaussian cdfs, which leads to an efficient method of approximating (41)-(42) in the case of *binary* logistic regression (i.e., $D = 2$) [38]. We now develop an extension of this method for the MLR case (i.e., $D \geq 2$).

To facilitate the Gaussian mixture (GM) approximation, we work with the difference variables

$$\gamma_d^{(y)} \triangleq \begin{cases} z_y - z_d & d \neq y \\ z_y & d = y \end{cases}. \quad (51)$$

Their utility can be seen from the fact that (recalling (3))

$$p_{y|z}(y|z) = \frac{1}{1 + \sum_{d \neq y} \exp(z_d - z_y)} \quad (52)$$

$$= \frac{1}{1 + \sum_{d \neq y} \exp(-\gamma_d^{(y)})} \triangleq l^{(y)}(\gamma^{(y)}), \quad (53)$$

which is smooth, positive, and bounded by 1, and strictly increasing in $\gamma_d^{(y)}$. Thus,⁴ for appropriately chosen $\{\alpha_l, \mu_{kl}, \sigma_{kl}\}$,

$$l^{(y)}(\gamma) \approx \sum_{l=1}^L \alpha_l \prod_{k \neq y} \Phi\left(\frac{\gamma_k - \mu_{kl}}{\sigma_{kl}}\right) \triangleq \hat{l}^{(y)}(\gamma), \quad (54)$$

where $\Phi(x)$ is the standard normal cdf, $\sigma_{kl} > 0$, $\alpha_l \geq 0$, and $\sum_l \alpha_l = 1$. In practice, the GM parameters $\{\alpha_l, \mu_{kl}, \sigma_{kl}\}$ could be designed off-line to minimize, e.g., the total variation distance $\sup_{\gamma \in \mathbb{R}^D} |l^{(y)}(\gamma) - \hat{l}^{(y)}(\gamma)|$.

Recall from (41)-(43) that our objective is to compute quantities of the form

$$\int_{\mathbb{R}^D} (e_d^\top z)^i p_{y|z}(y|z) \mathcal{N}(z; \hat{\mathbf{p}}, \mathbf{Q}^{\mathbf{p}}) dz \triangleq S_{di}^{(y)}, \quad (55)$$

where $i \in \{0, 1, 2\}$, $\mathbf{Q}^{\mathbf{p}}$ is diagonal, and e_d is the d th column of \mathbf{I}_D . To exploit (54), we change the integration variable to

$$\gamma^{(y)} = \mathbf{T}_y z \quad (56)$$

with

$$\mathbf{T}_y = \begin{bmatrix} -\mathbf{I}_{y-1} & \mathbf{1}_{(y-1) \times 1} & \mathbf{0}_{(y-1) \times (D-y)} \\ \mathbf{0}_{1 \times (y-1)} & 1 & \mathbf{0}_{1 \times (D-y)} \\ \mathbf{0}_{(D-y) \times (y-1)} & \mathbf{1}_{(D-y) \times 1} & -\mathbf{I}_{D-y} \end{bmatrix} \quad (57)$$

to get (since $\det(\mathbf{T}_y) = 1$)

$$S_{di}^{(y)} = \int_{\mathbb{R}^D} (e_d^\top \mathbf{T}_y^{-1} \gamma)^i l^{(y)}(\gamma) \mathcal{N}(\gamma; \mathbf{T}_y \hat{\mathbf{p}}, \mathbf{T}_y \mathbf{Q}^{\mathbf{p}} \mathbf{T}_y^\top) d\gamma. \quad (58)$$

Then, applying the approximation (54) and

$$\begin{aligned} \mathcal{N}(\gamma; \mathbf{T}_y \hat{\mathbf{p}}, \mathbf{T}_y \mathbf{Q}^{\mathbf{p}} \mathbf{T}_y^\top) &= \mathcal{N}(\gamma_y; \hat{p}_y, q_y^{\mathbf{p}}) \\ &\times \prod_{k \neq y} \mathcal{N}(\gamma_k; \gamma_y - \hat{p}_k, q_k^{\mathbf{p}}) \end{aligned} \quad (59)$$

to (58), we find that

$$\begin{aligned} S_{di}^{(y)} &\approx \sum_{l=1}^L \alpha_l \int_{\mathbb{R}} \mathcal{N}(\gamma_y; \hat{p}_y, q_y^{\mathbf{p}}) \left[\int_{\mathbb{R}^{D-1}} (e_d^\top \mathbf{T}_y^{-1} \gamma)^i \right. \\ &\times \left. \prod_{k \neq y} \mathcal{N}(\gamma_k; \gamma_y - \hat{p}_k, q_k^{\mathbf{p}}) \Phi\left(\frac{\gamma_k - \mu_{kl}}{\sigma_{kl}}\right) d\gamma_k \right] d\gamma_y. \end{aligned} \quad (60)$$

Noting that $\mathbf{T}_y^{-1} = \mathbf{T}_y$, we have

$$e_d^\top \mathbf{T}_y^{-1} \gamma = \begin{cases} \gamma_y - \gamma_d & d \neq y \\ \gamma_y & d = y \end{cases}. \quad (61)$$

⁴Note that, since the role of y in $\hat{l}^{(y)}(\gamma)$ is merely to ignore the y th component of the input γ , we could have instead written $\hat{l}^{(y)}(\gamma) = \hat{l}(\mathbf{J}_y \gamma)$ for y -invariant $\hat{l}(\cdot)$ and \mathbf{J}_y constructed by removing the y th row from the identity matrix.

Thus, for a fixed value of $\gamma_y = c$, the inner integral in (60) can be expressed as a product of linear combinations of terms

$$\int_{\mathbb{R}} \gamma^i \mathcal{N}(\gamma; c - \hat{p}, q) \Phi\left(\frac{\gamma - \mu}{\sigma}\right) d\gamma \triangleq T_i \quad (62)$$

with $i \in \{0, 1, 2\}$, which can be computed in closed form. In particular, defining $x \triangleq \frac{c - \hat{p} - \mu}{\sqrt{\sigma^2 + q}}$, we have

$$T_0 = \Phi(x) \quad (63)$$

$$T_1 = (c - \hat{p})\Phi(x) + \frac{q\phi(x)}{\sqrt{\sigma^2 + q}} \quad (64)$$

$$T_2 = \frac{(T_1)^2}{\Phi(x)} + q\Phi(x) - \frac{q^2\phi(x)}{\sigma^2 + q} \left(x + \frac{\phi(x)}{\Phi(x)}\right), \quad (65)$$

which can be obtained using the results in [39, §3.9]. The outer integral in (60) can then be approximated via numerical integration.

If a grid of K values is used for numerical integration over γ_y in (60), then the overall complexity of the method grows as $O(MDLK)$. Our experiments indicate that relatively small values (e.g., $L = 2$ and $K = 7$) suffice.

5) Performance comparison: Above we described four methods of approximating lines 15-16 in Algorithm 1 under diagonal $\mathbf{Q}^{\mathbf{p}}$ and $\mathbf{Q}^{\mathbf{z}}$. We now compare the accuracy and complexity of these methods. In particular, we measured the accuracy of the conditional mean (i.e., line 15) approximation as follows (for a given $\hat{\mathbf{p}}$ and $\mathbf{Q}^{\mathbf{p}}$):

- 1) generate i.i.d. samples $\mathbf{z}_{\text{true}}[t] \sim \mathcal{N}(\mathbf{z}; \hat{\mathbf{p}}, \mathbf{Q}^{\mathbf{p}})$ and $y_{\text{true}}[t] \sim p_{y|z}(y | \mathbf{z}_{\text{true}}[t])$ for $t = 1 \dots T$,
- 2) compute the approximation $\hat{\mathbf{z}}[t] \approx \mathbb{E}\{\mathbf{z} | y = y_{\text{true}}[t], \mathbf{p} = \hat{\mathbf{p}}; \mathbf{Q}^{\mathbf{p}}\}$ using each method described in Sections III-C1–III-C4,
- 3) compute average MSE $\triangleq \frac{1}{T} \sum_{t=1}^T \|\mathbf{z}_{\text{true}}[t] - \hat{\mathbf{z}}[t]\|_2^2$ for each method,

and we measured the combined runtime of lines 15-16 for each method. Unless otherwise noted, we used $D = 4$ classes, $\hat{\mathbf{p}} = \mathbf{e}_1$, $\mathbf{Q}^{\mathbf{p}} = q^{\mathbf{p}} \mathbf{I}_D$, and $q^{\mathbf{p}} = 1$ in our experiments. For numerical integration (NI), we used a grid of size $K = 7$ and radius of $\alpha = 4$ standard deviations; for importance sampling (IS), we used $K = 1500$ samples; and for the Gaussian-mixture (GM) method, we used $L = 2$ mixture components and a grid size of $K = 7$. Empirically, we found that smaller grids or fewer samples compromised accuracy, whereas larger grids or more samples compromised runtime.

Figure 2 plots the normalized MSE versus variance $q^{\mathbf{p}}$ for the four methods under test, in addition to the trivial method $\hat{\mathbf{z}}[t] = \hat{\mathbf{p}}$. The figure shows that the NI, IS, and GM methods performed similarly across the full range of $q^{\mathbf{p}}$ and always outperform the trivial method. The Taylor-series method, however, breaks down when $q^{\mathbf{p}} > 1$. A close examination of the figure reveals that GM gave the best accuracy, IS the second best accuracy, and NI the third best accuracy.

Figure 3 shows the cumulative runtime (over $M = 500$ training samples) of the methods from Sections III-C1–III-C4 versus the number of classes, D . Although the Taylor-series method was the fastest, we saw in Fig. 2 that it is accurate only at small variances $q^{\mathbf{p}}$. Figure 3 then shows GM was about an




Fig. 2: MSE/q^P versus variance q^P for various methods to compute line 15 in Algorithm 1. Each point represents the average of 5×10^6 independent trials.




Fig. 3: Cumulative runtime (over $M = 500$ samples) versus number-of-classes D for various methods to compute lines 15-16 in Algorithm 1. Each point represents the average of 2000 independent trials.

order-of-magnitude faster than IS, which was several orders-of-magnitude faster than NI.

Together, Figures 2-3, show that our proposed GM method dominated the IS and NI methods in both accuracy and runtime. Thus, for the remainder of the paper, we implement sum-product SHyGAMP using the GM method from Section III-C4.

D. Min-sum SHyGAMP: Inference of x_n

With diagonal \mathbf{Q}_n^r and \mathbf{Q}_n^x , the implementation of lines 3-4 in Algorithm 1 can be significantly simplified. Recall that, when the prior p_x is chosen as i.i.d. Laplace (5), line 3 manifests as (21), which is in general a non-trivial optimization problem. But with diagonal \mathbf{Q}_n^r , (21) decouples into D instances of the scalar optimization

$$x_{nd} = \arg \max_x -\frac{1}{2} \frac{(x - \hat{r}_{nd})^2}{q_{nd}^r} - \lambda|x|, \quad (66)$$

Algorithm	Quantity	Method	Complexity
SPA-SHyGAMP	\hat{x}	CF	$O(D)$
	\mathbf{Q}^x	CF	$O(D)$
	\hat{z}	GM	$O(LKD)$
	\mathbf{Q}^z	GM	$O(LKD)$
MSA-SHyGAMP	\hat{x}	ST	$O(D)$
	\mathbf{Q}^x	CF	$O(D)$
	\hat{z}	CWN	$O(KD)$
	\mathbf{Q}^z	CF	$O(D^3)$

TABLE II: A summary of the D -dimensional inference sub-problems encountered when running SPA-SHyGAMP or MSA-SHyGAMP, as well as their associated computational costs. ‘CF’ = ‘closed form’, ‘GM’ = ‘Gaussian mixture’, ‘ST’ = ‘Soft-thresholding’, and ‘CWN’ = ‘component-wise Newton’s method’. For the GM, L denotes the number of mixture components and K the number of samples in the 1D numerical integral, and for CWN K denotes the number of iterations.

which is known to have the closed-form “soft thresholding” solution

$$\hat{x}_{nd} = \text{sgn}(\hat{r}_{nd}) \max\{0, |\hat{r}_{nd}| - \lambda q_{nd}^r\}. \quad (67)$$

Above, $\text{sgn}(r) = 1$ when $r \geq 0$ and $\text{sgn}(r) = -1$ when $r < 0$.

Meanwhile, line 4 reduces to

$$q_{nd}^x = \left[\frac{\partial^2}{\partial x^2} \left(\frac{1}{2} \frac{(x - \hat{r}_{nd})^2}{q_{nd}^r} + \lambda|x| \right) \Big|_{x=\hat{x}_{nd}} \right]^{-1}, \quad (68)$$

which equals q_{nd}^r when $\hat{x}_{nd} \neq 0$ and is otherwise undefined. When $\hat{x}_{nd} = 0$, we set $q_{nd}^x = 0$.

E. Min-sum SHyGAMP: Inference of z_m

With diagonal \mathbf{Q}_m^p and \mathbf{Q}_m^z , the implementation of lines 12-13 in Algorithm 1 also simplifies. Recall that, when the likelihood $p_{y|z}$ takes the MLR form in (3), line 12 manifests as (26), which can be solved using a component-wise Newton’s method as in (31)-(33) for any \mathbf{Q}_m^p and \mathbf{Q}_m^z . When \mathbf{Q}_m^p is diagonal, the first and second derivatives (32)-(33) reduce to

$$g_{md}^{(k)} = p_{y|z}(d|\hat{z}_m^{(k)}) - \delta_{y_m-d} + (\hat{z}_{md}^{(k)} - \hat{p}_{md})/q_{md}^p. \quad (69)$$

$$H_{md}^{(k)} = p_{y|z}(d|\hat{z}_m^{(k)})^2 - p_{y|z}(d|\hat{z}_m^{(k)}) - 1/q_{md}^p, \quad (70)$$

which leads to a reduction in complexity.

Furthermore, line 13 simplifies, since with diagonal \mathbf{Q}_m^z it suffices to compute only the diagonal components of $\mathbf{H}_m^{(k)}$ in (30). In particular, when \mathbf{Q}_m^p is diagonal, the result becomes

$$q_{md}^z = \frac{1}{1/q_{md}^p + p_{y|z}(d|\hat{z}_m) - p_{y|z}(d|\hat{z}_m)^2}. \quad (71)$$

F. SHyGAMP summary

In summary, by approximating the covariance matrices as diagonal, the SPA-SHyGAMP and MSA-SHyGAMP algorithms improve computationally upon their HyGAMP counterparts. A summary of the D -dimensional inference problems encountered when running SPA-SHyGAMP or MSA-SHyGAMP, as well as their associated computational costs, is given in Table II. A high-level comparison between HyGAMP and SHyGAMP is given in Table III.

Algorithm	HyGAMP	SHyGAMP
Diagonal covariance matrices		✓
Simplified D -dimensional inference		✓
Scalar-variance approximation		✓
Online parameter tuning		✓

TABLE III: High-level comparison of SHyGAMP and HyGAMP.

IV. ONLINE PARAMETER TUNING

The weight vector priors in (5) and (39) depend on modeling parameters that, in practice, must be tuned. Although cross-validation (CV) is the customary approach to tuning such model parameters, it can be very computationally costly, since each parameter must be tested over a grid of hypothesized values and over multiple data folds. For example, K -fold cross-validation tuning of P parameters using G hypothesized values of each parameter requires the training and evaluation of KG^P classifiers.

A. Parameter selection for Sum-product SHyGAMP

For SPA-SHyGAMP, we propose to use the zero-mean Bernoulli-Gaussian prior in (39), which has parameters β_d , m_d , and v_d . Instead of CV, we use the EM-GM-AMP framework described in [25] to tune these parameters online. See [31] for details regarding the initialization of β_d , m_d , and v_d .

B. Parameter selection for Min-sum SHyGAMP

To use MSA-SHyGAMP with the Laplacian prior in (5), we need to specify the scale parameter λ . For this, we use a modification of the SURE-AMP framework from [26], which adjusts λ to minimize the Stein's unbiased risk estimate (SURE) of the weight-vector MSE.

We describe our method by first reviewing SURE and SURE-AMP. First, suppose that the goal is to estimate the value of x , which is a realization of the random variable \mathbf{x} , from the noisy observation r , which is a realization of

$$\mathbf{r} = \mathbf{x} + \sqrt{q^{\mathbf{r}}}\mathbf{w}, \quad (72)$$

with $\mathbf{w} \sim \mathcal{N}(0, 1)$ and $q^{\mathbf{r}} > 0$. For this purpose, consider an estimate of the form $\hat{x} = f(r, q^{\mathbf{r}}; \boldsymbol{\theta})$ where $\boldsymbol{\theta}$ contains tunable parameters. For convenience, define the shifted estimation function $g(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \triangleq f(r, q^{\mathbf{r}}; \boldsymbol{\theta}) - r$ and its derivative $g'(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \triangleq \frac{\partial}{\partial r} g(r, q^{\mathbf{r}}; \boldsymbol{\theta})$. Then Stein [40] established the following result on the mean-squared error, or risk, of the estimate \hat{x} :

$$\mathbb{E} \{ [\hat{x} - x]^2 \} = q^{\mathbf{r}} + \mathbb{E} \{ g^2(r, q^{\mathbf{r}}; \boldsymbol{\theta}) + 2q^{\mathbf{r}} g'(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \}. \quad (73)$$

The implication of (73) is that, given only the noisy observation r and the noise variance $q^{\mathbf{r}}$, one can compute an estimate

$$\text{SURE}(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \triangleq q^{\mathbf{r}} + g^2(r, q^{\mathbf{r}}; \boldsymbol{\theta}) + 2q^{\mathbf{r}} g'(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \quad (74)$$

of the $\text{MSE}(\boldsymbol{\theta}) \triangleq \mathbb{E} \{ [\hat{x} - x]^2 \}$ that is unbiased, i.e.,

$$\mathbb{E} \{ \text{SURE}(r, q^{\mathbf{r}}; \boldsymbol{\theta}) \} = \text{MSE}(\boldsymbol{\theta}). \quad (75)$$

These unbiased risk estimates can then be used as a surrogate for the true MSE when tuning $\boldsymbol{\theta}$.

In [26], it was noticed that the assumption (72) is satisfied by AMP's denoiser inputs $\{\hat{r}_n\}_{n=1}^N$, and thus [26] proposed to tune the soft threshold λ to minimize the SURE:

$$\hat{\lambda} = \arg \min_{\lambda} \sum_{n=1}^N g^2(\hat{r}_n, q^{\mathbf{r}}; \lambda) + 2q^{\mathbf{r}} g'(\hat{r}_n, q^{\mathbf{r}}; \lambda). \quad (76)$$

Recalling the form of the estimator $f(\cdot)$ from (67), we have

$$g^2(\hat{r}_n, q^{\mathbf{r}}; \lambda) = \begin{cases} \lambda^2 (q^{\mathbf{r}})^2 & \text{if } |\hat{r}_n| > \lambda q^{\mathbf{r}} \\ \hat{r}_n^2 & \text{otherwise} \end{cases} \quad (77)$$

$$g'(\hat{r}_n, q^{\mathbf{r}}; \lambda) = \begin{cases} -1 & \text{if } |\hat{r}_n| < \lambda q^{\mathbf{r}} \\ 0 & \text{otherwise} \end{cases}. \quad (78)$$

However, solving (76) for λ is non-trivial because the objective is non-smooth and has many local minima. A stochastic gradient descent approach was proposed in [26], but its convergence speed is too slow to be practical.

Since (72) also matches the scalar-variance SHyGAMP model from Section III-A, we propose to use SURE to tune λ for min-sum SHyGAMP. But, instead of the empirical average in (76), we propose to use a statistical average, i.e.,

$$\hat{\lambda} = \arg \min_{\lambda} \mathbb{E} \{ \underbrace{g^2(\mathbf{r}, q^{\mathbf{r}}; \lambda) + 2q^{\mathbf{r}} g'(\mathbf{r}, q^{\mathbf{r}}; \lambda)}_{\triangleq J(\lambda)} \}, \quad (79)$$

by modeling the random variable \mathbf{r} as a Gaussian mixture (GM) whose parameters are fitted to $\{\hat{r}_{nd}\}$. As a result, the objective in (79) is smooth. Moreover, by constraining the smallest mixture variance to be at least $q^{\mathbf{r}}$, the objective becomes unimodal, in which case $\hat{\lambda}$ from (79) is the unique root of $\frac{d}{d\lambda} J(\lambda)$. To find this root, we use the bisection method. In particular, due to (77)-(78), the objective in (79) becomes

$$J(\lambda) = \int_{-\infty}^{-\lambda q^{\mathbf{r}}} p_{\mathbf{r}}(r) \lambda^2 (q^{\mathbf{r}})^2 dr + \int_{-\lambda q^{\mathbf{r}}}^{\lambda q^{\mathbf{r}}} p_{\mathbf{r}}(r) (r^2 - 2q^{\mathbf{r}}) dr + \int_{\lambda q^{\mathbf{r}}}^{\infty} p_{\mathbf{r}}(r) \lambda^2 (q^{\mathbf{r}})^2 dr, \quad (80)$$

from which it can be shown that [31]

$$\begin{aligned} \frac{d}{d\lambda} J(\lambda) &= 2\lambda (q^{\mathbf{r}})^2 [1 - \Pr\{-\lambda q^{\mathbf{r}} < \mathbf{r} < \lambda q^{\mathbf{r}}\}] \\ &\quad - [p_{\mathbf{r}}(\lambda q^{\mathbf{r}}) + p_{\mathbf{r}}(-\lambda q^{\mathbf{r}})] 2(q^{\mathbf{r}})^2. \end{aligned} \quad (81)$$

For GM fitting, we use the standard EM approach [9] and find that relatively few (e.g., $L = 3$) mixture terms suffice. Note that we re-tune λ using the above technique at each iteration of Algorithm 1, immediately before line 3. Experimental verification of our method is provided in Section V-B.

V. NUMERICAL RESULTS

In this section we describe the results of several experiments used to test SHyGAMP. In these experiments, EM-tuned SPA-SHyGAMP and SURE-tuned MSA-SHyGAMP were compared to two state-of-the-art sparse MLR algorithms: SBMLR [14] and GLMNET [13]. We are particularly interested in SBMLR and GLMNET because [13,14] show that they have strong advantages over earlier algorithms, e.g., [10,11,12]. As described in Section I-C, both SBMLR and GLMNET

use ℓ_1 regularization, but SBMLR tunes the regularization parameter λ using evidence maximization while GLMNET tunes it using cross-validation (using the default value of 10 folds unless otherwise noted). For SBMLR and GLMNET, we ran code written by the authors⁵⁶ under default settings (unless otherwise noted). For SHyGAMP, we used the damping modification described in [30]. We note that the runtimes reported for all algorithms include the total time spent to tune all parameters and train the final classifier.

Due to space limitations, we do not show the performance of the more complicated HyGAMP algorithm from Section II. However, our experience suggests that HyGAMP generates weight matrices $\hat{\mathbf{X}}$ that are very similar to those generated by SHyGAMP, but with much longer runtimes, especially as D grows.

A. Synthetic data in the $M \ll N$ regime

We first describe the results of three experiments with synthetic data. For these experiments, the training data were randomly generated and algorithm performance was averaged over several data realizations. In all cases, we started with balanced training labels $y_m \in \{1, \dots, D\}$ for $m = 1, \dots, M$ (i.e., M/D examples from each of D classes). Then, for each data realization, we generated M i.i.d. training features \mathbf{a}_m from the class-conditional generative distribution $\mathbf{a}_m | y_m \sim \mathcal{N}(\boldsymbol{\mu}_{y_m}, v\mathbf{I}_N)$. In doing so, we chose the intra-class variance, v , to attain a desired Bayes error rate (BER) of 10% (see [31] for details), and we used randomly generated K -sparse orthonormal class means, $\boldsymbol{\mu}_d \in \mathbb{R}^N$. In particular, we generated $[\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_D]$ by drawing a $K \times K$ matrix with i.i.d. $\mathcal{N}(0, 1)$ entries, performing a singular value decomposition, and zero-padding the first D left singular vectors to length N . We note that our generation of $\mathbf{y}, \mathbf{A}, \mathbf{X}$ is matched [41] to the multinomial logistic model (2)-(3).

Given a training data realization, each algorithm was invoked to yield a weight matrix $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_D]$. The corresponding expected test-error rate was then analytically computed as

$$\Pr\{\text{err}\} = 1 - \frac{1}{D} \sum_{y=1}^D \Pr\{\text{cor}|y\} \quad (82)$$

$$\Pr\{\text{cor}|y\} = \Pr \bigcap_{d \neq y} \left\{ (\hat{\mathbf{x}}_y - \hat{\mathbf{x}}_d)^\top \mathbf{a} < (\hat{\mathbf{x}}_y - \hat{\mathbf{x}}_d)^\top \boldsymbol{\mu}_y \right\}, \quad (83)$$

where $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, v\mathbf{I}_N)$ and the multivariate normal cdf in (83) was computed using Matlab's `mvncdf`.

For all three synthetic-data experiments, we used $D = 4$ classes and $K \ll M \ll N$. In the first experiment, we fixed K and N and we varied M ; in the second experiment, we fixed K and M and we varied K ; and in the third experiment, we fixed K and M and we varied N . The specific values/ranges of K, M, N used for each experiment are given in Table IV.

Figures 4a-b show the expected test-error rate and runtime, respectively, versus the number of training examples, M , averaged over 12 independent trials. Figure 4a shows that,

Experiment	M	N	K	D
1	$\{100, \dots, 5000\}$	10000	10	4
2	300	30000	$\{5, \dots, 30\}$	4
3	200	$\{10^3, \dots, 10^{5.5}\}$	10	4
4	300	30000	25	4

TABLE IV: Configurations of the synthetic-data experiments.

at all tested values of M , SPA-SHyGAMP gave the best error-rates and MSA-SHyGAMP gave the second best error-rates, although those reached by GLMNET were similar at large M . Moreover, the error-rates of SPA-SHyGAMP, MSA-SHyGAMP, and GLMNET all converged towards the BER as M increased, whereas that of SBMLR did not. Since MSA-SHyGAMP, GLMNET, and SBMLR all solve the same ℓ_1 -regularized MLR problem, the difference in their error-rates can be attributed to the difference in their tuning of the regularization parameter λ . Figure 4b shows that, for $M > 500$, SPA-SHyGAMP was the fastest, followed by MSA-SHyGAMP, SBMLR, and GLMNET. Note that the runtimes of SPA-SHyGAMP, MSA-SHyGAMP, and GLMNET increased linearly with M , whereas the runtime of SBMLR increased quadratically with M .

Figures 5a-b show the expected test-error rate and runtime, respectively, versus feature-vector sparsity, K , averaged over 12 independent trials. Figure 5a shows that, at all tested values of K , SPA-SHyGAMP gave the best error-rates and MSA-SHyGAMP gave the second best error-rates. Figure 5b shows that SPA-SHyGAMP and MSA-SHyGAMP gave the fastest runtimes. All runtimes were approximately invariant to K .


Figures 6a-b show the expected test-error rate and runtime, respectively, versus the number of features, N , averaged over 12 independent trials. Figure 6a shows that, at all tested values of N , MSA-SHyGAMP gave lower error-rates than SBMLR and GLMNET. Meanwhile, SPA-SHyGAMP gave the lowest error-rates for certain values of N . Figure 6b shows that SPA-SHyGAMP and MSA-SHyGAMP gave the fastest runtimes for $N \geq 10000$, while SBMLR gave the fastest runtimes for $N \leq 3000$. All runtimes increased linearly with N .

B. Example of SURE tuning


Although the good error-rate performance of MSA-SHyGAMP in Section V-A suggests that the SURE λ -tuning method from Section IV-B is working reliably, we now describe a more direct test of its behavior. Using synthetic data generated as described in Section V-A with $D = 4$ classes, $N = 30000$ features, $M = 300$ examples, and sparsity $K = 25$, we ran MSA-SHyGAMP using various fixed values of λ . In the sequel, we refer to this experiment as ‘‘Synthetic Experiment 4.’’ The resulting expected test-error rate versus λ (averaged over 10 independent realizations) is shown in Fig. 7. For the same realizations, we ran MSE-SHyGAMP with SURE-tuning and plot the resulting error-rate and average $\hat{\lambda}$ in Fig. 7. From Fig. 7, we see that the SURE λ -tuning method matched both the minimizer and the minimum of the error-versus- λ trace of fixed- λ MSA-SHyGAMP.

⁵SBMLR obtained from <http://theoval.cmp.uea.ac.uk/matlab/>

⁶GLMNET obtained from http://www.stanford.edu/~hastie/glmnet_matlab/




(a) Error




(b) Runtime

Fig. 4: Synthetic Experiment 1: expected test-error rate and runtime versus M . Here, $D = 4$, $N = 10000$, and $K = 10$.



(a) Error



(b) Runtime

Fig. 5: Synthetic Experiment 2: expected test-error rate and runtime versus K . Here, $D = 4$, $M = 300$, and $N = 30000$.

C. Micro-array gene expression


Next we consider classification and feature-selection using micro-array gene expression data. Here, the labels indicate which type of disease is present (or no disease) and the features represent gene expression levels. The objective is i) to determine which subset of genes best predicts the various diseases and ii) to classify whether an (undiagnosed) patient is at risk for any of these diseases based on their gene profile.

We tried two datasets: one from Sun et al. [1] and one from Bhattacharjee et al. [2]. The Sun dataset includes $M = 179$ examples, $N = 54613$ features, and $D = 4$ classes; and the Bhattacharjee dataset includes $M = 203$ examples, $N = 12600$ features, and $D = 5$ classes. With the Sun dataset, we applied a $\log_2(\cdot)$ transformation and z-scored prior to processing, while with Bhattacharjee we simply z-scored (since the dataset included negative values).


The test-error rate was estimated as follows for each dataset. We consider a total of T “trials.” For the t th trial, we i) partition the dataset into a training subset of size $M_{\text{train},t}$ and

a test subset of size $M_{\text{test},t}$, ii) design the classifier using the training subset, and iii) apply the classifier to the test subset, recording the test errors $\{e_{tm}\}_{m=1}^{M_{\text{test},t}}$, where $e_{tm} \in \{0,1\}$ indicates whether the m th example was in error. We then estimate the average test-error rate using the empirical average $\hat{\mu} \triangleq M_{\text{test}}^{-1} \sum_{t=1}^T \sum_{m=1}^{M_{\text{test},t}} e_{tm}$, where $M_{\text{test}} = \sum_{t=1}^T M_{\text{test},t}$. If the test sets are constructed without overlap, we can model $\{e_{tm}\}$ as i.i.d. Bernoulli(μ), where μ denotes the true test-error rate. Then, since $\hat{\mu}$ is Binomial(μ, M_{test}), the standard deviation (SD) of our error-rate estimate $\hat{\mu}$ is $\sqrt{\text{var}\{\hat{\mu}\}} = \sqrt{\mu(1-\mu)/M_{\text{test}}}$. Since μ is unknown, we approximate the SD by $\sqrt{\hat{\mu}(1-\hat{\mu})/M_{\text{test}}}$.

Tables V and VI show, for each algorithm, the test-error rate estimate $\hat{\mu}$, the approximate SD $\sqrt{\hat{\mu}(1-\hat{\mu})/M_{\text{test}}}$ of the estimate, the average runtime, and two metrics for the sparsity of $\hat{\mathbf{X}}$. The $\|\hat{\mathbf{X}}\|_0$ metric quantifies the number of non-zero entries in $\hat{\mathbf{X}}$ (i.e., absolute sparsity), while the \hat{K}_{99} metric quantifies the number of entries of $\hat{\mathbf{X}}$ needed to reach 99% of the Frobenius norm of $\hat{\mathbf{X}}$ (i.e., effective sparsity). We note that



(a) Error



(b) Runtime

Fig. 6: Synthetic Experiment 3: expected test-error rate and runtime versus N . Here, $D = 4$, $M = 200$, and $K = 10$.

the reported values of \hat{K}_{99} and $\|\hat{\mathbf{X}}\|_0$ represent the average over the T folds. For both the Sun and Bhattacharjee datasets, we used $T = 19$ trials and $M_{\text{test},t} = \lfloor M/20 \rfloor \forall t$.

Table V shows results for the Sun dataset. There we see that MSA-SHyGAMP gave the best test-error rate, although the other algorithms were not far behind and all error-rate estimates were within the estimator standard deviation. SPA-SHyGAMP was the fastest algorithm and MSA-SHyGAMP was the second fastest, with the remaining algorithms running $3\times$ to $5\times$ slower than SPA-SHyGAMP. GLMNET's weights were the sparsest according to both sparsity metrics. SPA-SHyGAMP's weights had the second lowest value of \hat{K}_{99} , even though they were technically non-sparse (i.e., $\|\hat{\mathbf{X}}\|_0 = 218452 = ND$) as expected. Meanwhile, MSA-SHyGAMP's weights were the least sparse according to the \hat{K}_{99} metric.

Table VI shows results for the Bhattacharjee dataset. In this experiment, SPA-SHyGAMP and SBMLR were tied for the best error rate, MSA-SHyGAMP was 0.5 standard-deviations worse, and GLMNET was 1.2 standard-deviations worse.




Fig. 7: Synthetic experiment 4: expected test-error rate versus regularization parameter λ for fixed- λ MSA-SHyGAMP. Here, $D = 4$, $M = 300$, $N = 30000$, and $K = 25$. Also shown is the average test-error rate for SURE-tuned MSA-SHyGAMP plotted at the average value of $\hat{\lambda}$.

Algorithm	% Error (SD)	Runtime (s)	\hat{K}_{99}	$\ \hat{\mathbf{X}}\ _0$
SPA-SHyGAMP	33.3 (3.8)	6.86	20.05	218452
MSA-SHyGAMP	31.0 (3.7)	13.59	93.00	145.32
SBMLR	31.6 (3.7)	22.48	49.89	72.89
GLMNET	33.9 (3.8)	31.93	10.89	16.84

TABLE V: Estimated test-error rate, standard deviation of estimate, runtime, and sparsities for the Sun dataset.

However, SPA-SHyGAMP ran about twice as fast as SBMLR, and $4\times$ as fast as GLMNET. As in the Sun dataset, SPA-SHyGAMP returned the sparsest weight matrix according to the \hat{K}_{99} metric. The sparsities of the weight matrices returned by the other three algorithms were similar to one another in both metrics. Unlike in the Sun dataset, MSA-SHyGAMP and SBMLR had similar runtimes (which is consistent with Fig. 6b since N is lower here than in the Sun dataset).

D. Text classification with the RCV1 dataset

Next we consider text classification using the Reuter's Corpus Volume 1 (RCV1) dataset [6]. Here, each sample (y_m, \mathbf{a}_m) represents a news article, where y_m indicates the article's topic and \mathbf{a}_m indicates the frequencies of common words in the article. The version of the dataset that we used⁷ contained $N = 47236$ features and 53 topics. However, we used only the first $D = 25$ of these topics (to reduce the computational demand). Also, we retained the default training and test partitions, which resulted in the use of $M = 14147$ samples for training and 469571 samples for testing.

The RCV1 features are very sparse (only 0.326% of the features are non-zero) and non-negative, which conflicts with the standard assumptions used for the derivation of AMP algorithms: that \mathbf{A} is i.i.d. zero-mean and sub-Gaussian. Interestingly, the RCV1 dataset also caused difficulties for SBMLR, which diverged under default settings. This divergence was

⁷<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multiclass.html>

Algorithm	% Error (SD)	Runtime (s)	\hat{K}_{99}	$\ \hat{\mathbf{X}}\ _0$
SPA-SHyGAMP	9.5 (2.1)	3.26	16.15	63 000
MSA-SHyGAMP	10.5 (2.2)	6.11	55.20	84.65
SBMLR	9.5 (2.1)	6.65	44.25	79.10
GLMNET	12.0 (2.4)	13.67	49.65	89.40

TABLE VI: Estimated test-error rate, standard deviation of estimate, runtime, and sparsities for the Bhattacharjee dataset.




Fig. 8: Test-error rate versus runtime for the RCV1 dataset.

remedied by decreasing the value of a step-size parameter⁸ to 0.1 from the default value of 1.

Figure 8 shows test-error rate versus runtime for SPA-SHyGAMP, MSA-SHyGAMP, SBMLR, and GLMNET on the RCV1 dataset. In the case of SPA-SHyGAMP, MSA-SHyGAMP and SBMLR, each point in the figure represents one iteration of the corresponding algorithm. For GLMNET, each data-point represents one iteration of the algorithm *after* its cross-validation stage has completed.⁹ We used 2 CV folds (rather than the default 10) in this experiment to avoid excessively long runtimes. The figure shows that the SHyGAMP algorithms converged more than an order-of-magnitude faster than SBMLR and GLMNET, although the final error rates were similar. SPA-SHyGAMP displayed faster initial convergence, but MSA-SHyGAMP eventually caught up.

E. MNIST handwritten digit recognition

Finally, we consider handwritten digit recognition using the Mixed National Institute of Standards and Technology (MNIST) dataset [42]. This dataset consists of 70 000 examples, where each example is an $N = 784$ pixel image of one of $D = 10$ digits between 0 and 9. These features were again non-negative, which conflicts with the standard AMP assumption of i.i.d. zero-mean \mathbf{A} .

Our experiment characterized test-error rate versus the number of training examples, M , for the SPA-SHyGAMP, MSA-SHyGAMP, SBMLR, and GLMNET algorithms. For each

⁸See the variable `scale` on lines 129 and 143 of `sbmlr.m`.

⁹GLMNET spent most of its time on cross-validation. After cross-validation, GLMNET took 25.26 seconds to run, which is similar to the total runtimes of SPA-SHyGAMP and MSE-SHyGAMP.




Fig. 9: Estimated test-error rate versus M for the MNIST dataset, with error-bars indicating the standard deviation of the estimate.

value of M , we performed 50 Monte-Carlo trials. In each trial, M training samples were selected uniformly at random and the remainder of the data were used for testing. Figure 9 shows the average estimated test-error rate $\hat{\mu}$ versus the number of training samples, M , for the algorithms under test. The error-bars in the figure correspond to the average of the per-trial estimated SD over the 50 trials. For SBMLR, we reduced the stepsize to 0.5 from the default value of 1 to prevent a significant degradation of test-error rate. The figure shows SPA-SHyGAMP attaining significantly better error-rates than the other algorithms at small values of M (and again at the largest value of M considered for the plot). For this plot, M was chosen to focus on the $M < N$ regime.

VI. CONCLUSION

For the problem of multi-class linear classification and feature selection, we proposed several AMP-based approaches to sparse multinomial logistic regression. We started by proposing two algorithms based on HyGAMP [18], one of which finds the maximum a posteriori (MAP) linear classifier based on the multinomial logistic likelihood and a Laplacian prior, and the other of which finds an approximation of the test-error-rate minimizing linear classifier based on the multinomial logistic likelihood and a Bernoulli-Gaussian prior. The numerical implementation of these algorithms is challenged, however, by the need to solve D -dimensional inference problems of multiplicity M at each HyGAMP iteration. Thus, we proposed simplified HyGAMP (SHyGAMP) approximations based on a diagonalization of the message covariances and a careful treatment of the D -dimensional inference problems. In addition, we described EM- and SURE-based methods to tune the hyperparameters of the assumed statistical model. Finally, using both synthetic and real-world datasets, we demonstrated improved error-rate and runtime performance relative to the state-of-the-art SBMLR [13] and GLMNET [14] approaches to sparse MLR.

REFERENCES

- [1] H. Sun *et al.*, “Neuronal and glioma-derived stem cell factor induces angiogenesis within the brain,” *Cancer Cell*, vol. 9, pp. 287–300, 2006.
- [2] A. Bhattacharjee *et al.*, “Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses,” *Proc. Nat. Acad. Sci.*, vol. 98, pp. 13790–13795, Nov. 2001.
- [3] J. Haxby, M. Gobbini, M. Furey, A. Ishai, J. Schouten, and P. Pietrini, “Distributed and overlapping representations of faces and objects in ventral temporal cortex,” *Science*, vol. 293, pp. 2425–2430, 2001.
- [4] K. A. Norman, S. M. Polyn, G. J. Detre, and J. V. Haxby, “Beyond mind-reading: multi-voxel pattern analysis of fMRI data,” *Trends in Cognitive Sciences*, vol. 10, pp. 424–430, Sep. 2006.
- [5] G. Forman, “An extensive empirical study of feature selection metrics for text classification,” *J. Mach. Learn. Res.*, vol. 3, pp. 1289–1305, 2003.
- [6] D. Lewis, Y. Yang, T. Rose, and F. Li, “RCV1: A new benchmark collection for text categorization research,” *Journal of Machine Learning Research*, vol. 5, pp. 361–397, April 2004.
- [7] A. Gustafsson, A. Hermann, and F. Huber, *Conjoint Measurement: Methods and Applications*. Berlin: Springer-Verlag, 2007.
- [8] Y. Plan and R. Vershynin, “Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach,” *IEEE Trans. Inform. Theory*, vol. 59, no. 1, pp. 482–494, 2013.
- [9] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2007.
- [10] M. E. Tipping, “Sparse Bayesian learning and the relevance vector machine,” *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, 2001.
- [11] B. Krishnapuram, L. Carin, M. Figueiredo, and A. Hartemink, “Sparse multinomial logistic regression: Fast algorithms and generalization bounds,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, pp. 957–968, June 2005.
- [12] A. Genkin, D. D. Lewis, and D. Madigan, “Large-scale Bayesian logistic regression for text categorization,” *Technometrics*, vol. 49, pp. 291–304, Aug. 2007.
- [13] J. Friedman, T. Hastie, and R. Tibshirani, “Regularization paths for generalized linear models via coordinate descent,” *J. Statist. Softw.*, vol. 33, pp. 1–22, Jan. 2010.
- [14] G. C. Cawley, N. L. C. Talbot, and M. Girolami, “Sparse multinomial logistic regression via Bayesian L1 regularisation,” in *Proc. Neural Inform. Process. Syst. Conf.*, pp. 209–216, 2007.
- [15] L. Meier, S. van de Geer, and P. Bühlmann, “The group lasso for logistic regression,” *J. Roy. Statist. Soc. B*, vol. 70, pp. 53–71, 2008.
- [16] Y. Grandvalet, “Least absolute shrinkage is equivalent to quadratic penalization,” in *Proc. Int. Conf. Artif. Neural Netw.*, pp. 201–206, 1998.
- [17] D. J. C. MacKay, “The evidence framework applied to classification networks,” *Neural Comput.*, vol. 4, pp. 720–736, 1992.
- [18] S. Rangan, A. K. Fletcher, V. K. Goyal, and P. Schniter, “Hybrid generalized approximate message passing with applications to structured sparsity,” in *Proc. IEEE Int. Symp. Inform. Thy.*, pp. 1236–1240, July 2012. (full version at *arXiv:1111.2581*).
- [19] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman, 1988.
- [20] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” in *Proc. IEEE Int. Symp. Inform. Thy.*, pp. 2168–2172, Aug. 2011. (full version at *arXiv:1010.5141*).
- [21] J. Ziniel, P. Schniter, and P. Sederberg, “Binary classification and feature selection via generalized approximate message passing,” *IEEE Trans. Signal Process.*, vol. 63, no. 8, pp. 2020–2032, 2015.
- [22] D. L. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing,” *Proc. Nat. Acad. Sci.*, vol. 106, pp. 18914–18919, Nov. 2009.
- [23] D. L. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing: I. Motivation and construction,” in *Proc. Inform. Theory Workshop*, (Cairo, Egypt), pp. 1–5, Jan. 2010.
- [24] D. A. Knowles and T. P. Minka, “Non-conjugate variational message passing for multinomial and binary regression,” in *Proc. Neural Inform. Process. Syst. Conf.*, 2011.
- [25] J. P. Vila and P. Schniter, “Expectation-maximization Gaussian-mixture approximate message passing,” *IEEE Trans. Signal Process.*, vol. 61, pp. 4658–4672, Oct. 2013.
- [26] A. Mousavi, A. Maleki, and R. G. Baraniuk, “Parameterless, optimal approximate message passing,” *arXiv:1311.0035*, Nov. 2013.
- [27] G. F. Cooper, “The computational complexity of probabilistic inference using Bayesian belief networks,” *Artificial Intelligence*, vol. 42, pp. 393–405, 1990.
- [28] A. Javanmard and A. Montanari, “State evolution for general approximate message passing algorithms, with applications to spatial coupling,” *Inform. Inference*, vol. 2, no. 2, pp. 115–144, 2013.
- [29] S. Rangan, P. Schniter, E. Riegler, A. Fletcher, and V. Cevher, “Fixed points of generalized approximate message passing with arbitrary matrices,” in *Proc. IEEE Int. Symp. Inform. Thy.*, pp. 664–668, July 2013. (full version at *arXiv:1301.6295*).
- [30] S. Rangan, P. Schniter, and A. Fletcher, “On the convergence of generalized approximate message passing with arbitrary matrices,” in *Proc. IEEE Int. Symp. Inform. Thy.*, pp. 236–240, July 2014. (full version at *arXiv:1402.3210*).
- [31] E. M. Byrne, “Sparse multinomial logistic regression via approximate message passing,” Master’s thesis, The Ohio State University, Columbus, Ohio, July 2015.
- [32] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. Roy. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [33] D. R. Hunter and K. Lange, “A tutorial on MM algorithms,” *The American Statistician*, vol. 58, no. 1, pp. 30–37, 2004.
- [34] D. Hunter and R. Li, “Variable selection using MM algorithms,” *Ann. Statist.*, vol. 33, no. 4, pp. 1617–1642, 2005.
- [35] D. Bertsekas, *Nonlinear Programming*. Athena Scientific, 2nd ed., 1999.
- [36] D. Böhning, “Multinomial logistic regression algorithm,” *Ann. Inst. Statist. Math.*, vol. 44, pp. 197–200, 1992.
- [37] P. Schniter, “Turbo reconstruction of structured sparse signals,” in *Proc. Conf. Inform. Science & Syst.*, (Princeton, NJ), pp. 1–6, Mar. 2010.
- [38] L. A. Stefanski, “A normal scale mixture representation of the logistic distribution,” *Stats. Prob. Letters*, vol. 11, no. 1, pp. 69–70, 1991.
- [39] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.
- [40] C. M. Stein, “Estimation of the mean of a multivariate normal distribution,” *Ann. Statist.*, vol. 9, pp. 1135–1151, 1981.
- [41] M. I. Jordan, “Why the logistic function? A tutorial discussion on probabilities and neural networks,” Tech. Rep. 9503, MIT, Computational Cognitive Science, 1995.
- [42] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” in *Proc. IEEE*, vol. 86, pp. 2278–2324, Nov. 1998.